Share based fairness for arbitrary entitlements

Uriel Feige

Weizmann Institute

The setting

A set M of m indivisible items (goods).

n agents.

Each agent i has a monotone valuation function v_i.

- Special case: additive valuation. $v_i(S) = \sum_{e \in S} v_i(e)$.

Each agent i has an arbitrary non-negative entitlements b_i.

- Special case: equal entitlements. $b_i = \frac{1}{n}$ for every agent i.

Assumptions about entitlements

Monotone: higher entitlement is better.

Linear scale: if agent j transfers her entitlement to agent i, then agent i has entitlement $b_i + b_j$.

Entitlements sum up to 1, indicating that collectively all agents are entitled to all the goods, no less and no more.

Examples

- Allocation of housing units to eligible residents.
- Allocation of seats in the parliament to political parties.
- Dividing an inheritance among a widow and several children.
- Handling student registration to courses of limited capacity.
- NBA draft: allocation of eligible basketball players to NBA teams.

…

Recap for equal entitlement

Ideal fairness notions: represent what we aspire to achieve.

Not always feasible, so sometimes we settle for approximations.

Share based notions:

- Maximin share (MMS) for monotone valuations.
- For additive valuations: the proportional share (PS).

Comparison based notions: envy-free.

Sanity check for fairness notions

Consider dividing a homogeneous divisible good.

Money, water, processing time, storage space ...

Modelled as m identical items, for very large m. $v_i(S) = v_i(|S|)$.

Proportional allocation: each agent i gets $b_i \cdot m$ items (m is assumed to be such that $b_i \cdot m$ is an integer).

Independent of the valuations.

In the equal entitlement case, the ideal fairness notions (MMS, PS for additive, EF) are satisfied by the proportional allocation.

Unequal entitlement, additive valuations

A set M of m identical items, $v_i(S) = |S|$.

Maximin share?

Proportional share (PS): entitled to a value of $b_i \cdot v_i(M)$.

Weighted envy free (WEF): $\frac{v_i(A_i)}{b_i} \geq \frac{v_i(A_j)}{b_j}$.

Both enforce the proportional allocation.

Non-additive valuations, share based notions

Example, submodular valuation (diminishing returns).

$$v_i(x) = \sqrt{x}$$

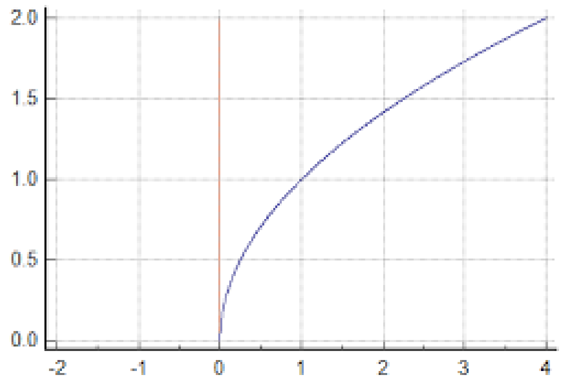

PS. Not meant to be used for non-additive valuations.

$PS_i = b_i \cdot v_i(M) = b_i \cdot \sqrt{m}$. Achieved by $(b_i)^2 \cdot m$ items.

The proportional share for nonadditive valuations does not give the proportional allocation, not even for equal entitlement.

How does one adapt MMS to unequal entitlements?

Non-additive valuations, WEF $\frac{v_i(A_i)}{b_i} \geq \frac{v_i(A_j)}{b_j}$

With two agents of equal entitlement, each agent gets $\frac{m}{2}$ items.

Unequal entitlements $\left(\frac{1}{3}, \frac{2}{3}\right)$, identical valuations $v_i(x) = \sqrt{x}$.

In WEF, agent 1 gets $\frac{m}{5}$ items and agent 2 gets $\frac{4m}{5}$ items.

WEF allocation offers lower Nash Social Welfare $v_1(A_1) \cdot v_2(A_2)$ and lower weighted NSW $(v_1(A_1))^{b_1}(v_2(A_2))^{b_2}$ then the proportional allocation.

WEF allocations do not exist if $v_1(x) = x$ and $v_2(x) = \sqrt{x}$.

WEF does not appear to capture fair allocations.

Ex-ante fairness (expected value received)

Giving M to each agent i with probability b_i gives the proportional share ex-ante, and is ex-ante WEF.

Weaknesses of this randomized allocation:

- Ex-post very poor. Want a combination of ex-ante and ex-post guarantees (Best of Both Worlds, BoBW).
- The ex-ante proportional guarantee is optimal for additive valuations, but might be very weak for concave (subadditive) valuations, in which the sum of values of the parts is higher than the value of the whole.

Definitions presented in this talk

Want fairness notions that apply to arbitrary (unequal) entitlements, and arbitrary monotone valuations.

Leave comparison based notions (envy-free) for future work.

Concentrate on share based fairness:

- $\widehat{MMS}$. A minimal upper bound on all feasible ex-post shares.
- $\widehat{MES}$. A minimal upper bound on all feasible ex-ante shares.
- APS. Gives proportional allocation for homogeneous divisible good.

Ex-post share-based fairness notions

Agent i has entitlement $b_i \geq 0$, and $\sum_i b_i = 1$.
(For equal entitlement, $b_i = \frac{1}{n}$.)

If $b_i = \frac{1}{k}$ for some integer k, it is natural to use the share value:

$$MMS\left(v_i, \frac{1}{k}\right) = \max_{\{A_1, \dots, A_k\}} \min_j v_i(A_j)$$

Suppose we do that.

What should the share value be when $\frac{1}{k+1} < b_i < \frac{1}{k}$?

Pessimistic rounding $\frac{1}{k+1} \leq b_i < \frac{1}{k}$

$$\widetilde{MMS}(v_i, b_i) = MMS(v_i, \frac{1}{k+1})$$

Possible motivation: feasibility when all agents have the same valuation function.

Proposition [Feige23]: there are instances with entitlements $(\frac{1}{2}, \frac{1}{3}, \frac{1}{6})$ and identical additive valuations with no $\widetilde{MMS}$ allocations.

Optimistic rounding

$$\frac{1}{k+1} < b_i \leq \frac{1}{k}$$

$$\widehat{MMS}(v_i, b_i) = MMS(v_i, \frac{1}{k})$$

Motivation: an upper bound on all feasible shares.

Proof: There might be k agents with entitlement b_i and valuation v_i, and then at least one of the agents will get at most $\widehat{MMS}(v_i, b_i)$.

This upper bound is tight, even with respect to nice feasible shares.

Theorem [Babaioff and Feige, 24]: there is a collection of shares, such that each share is feasible, *item name independent* and *self-maximizing*, and for every b_i and additive v_i, the value of at least one of these shares is $\widehat{MMS}(v_i, b_i)$.

The AnyPrice Share APS [Babaioff, Ezra, Feige 2021]

Agent i has entitlement $b_i \geq 0$, and $\sum_i b_i = 1$.
(For equal entitlement, $b_i = \frac{1}{n}$.)

Recall: the MMS of an agent is the value that she is guaranteed to get as a cutter in a *cut and choose* game.

At a high level: the APS of an agent is the value that she is guaranteed to get as a chooser in a *price and choose* game.

Price and choose can handle arbitrary entitlements.

If $b_i = \frac{1}{k}$ for some integer k, the APS need not equal $MMS\left(v_i, \frac{1}{k}\right)$.

The anyprice share APS

[Babaioff, Ezra, Feige 2021]

An adversary assigns a nonnegative vector $p = (p_1, \dots, p_m)$ of prices to the items, summing to 1.

The agent may pick any bundle B of price at most b_i.

$$APS = min_p max_B [v(B)]$$

A min-max definition.

The APS of agent i depends only on the valuation function and entitlement of i.

Additive valuation, entitlement $b_i = \frac{2}{5} = 0.4$

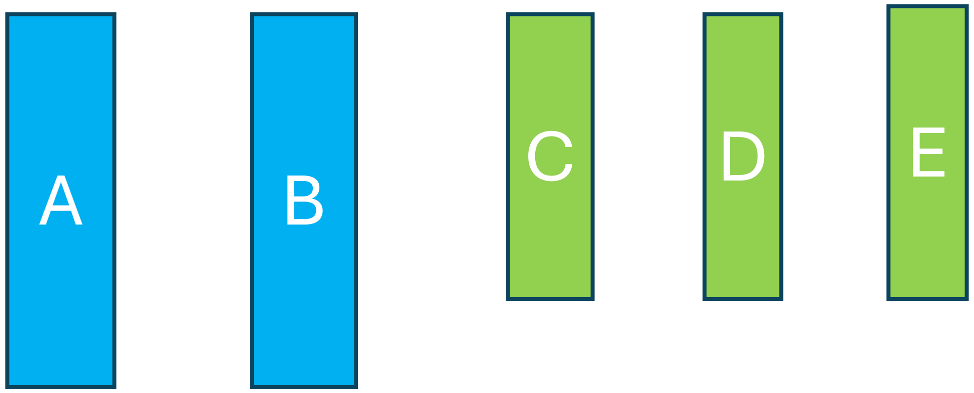

$$v_i(A) = v_i(B) = 7$$

$$v_i(C) = v_i(D) = v_i(E) = 5$$

Additive valuation, entitlement $b_i = \frac{2}{5} = 0.4$

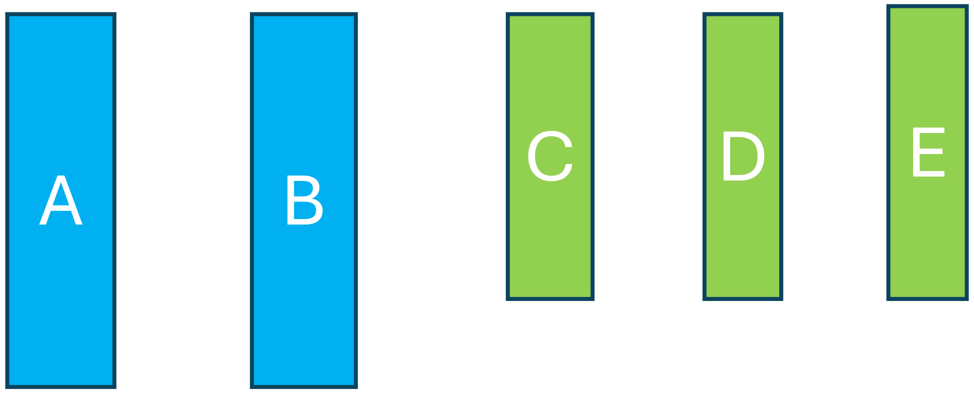

0.23 0.23 0.18 0.18 0.18

$$v_i(A) = v_i(B) = 7$$

$$v_i(C) = v_i(D) = v_i(E) = 5$$

$$APS_i \leq \max[5+5,\ 7] = 10$$

Additive valuation, entitlement $b_i = \frac{2}{5} = 0.4$

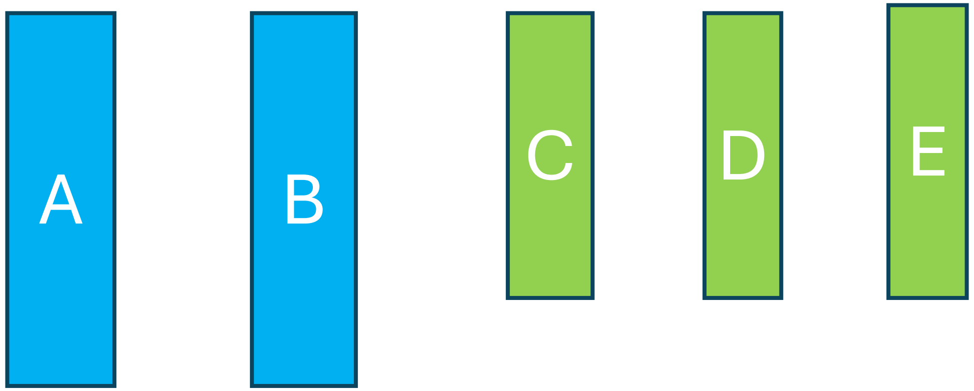

0.23 0.23 0.18 0.18 0.18

$v_i(A) = v_i(B) = 7$

$v_i(C) = v_i(D) = v_i(E) = 5$

$APS_i \leq \max[5+5,\ 7] = 10$

5 bundles, each item in two of them: $(A,B),(B,C),(C,D),(D,E),(E,A)$.

At least one bundle has price at most 0.4. Implies that $APS_i \geq 10$

Properties of APS

A distribution over subsets of $\mathcal{M}$ is b_i-balanced if every item has probability b_i of appearing in a random bundle.

Max-min definition of APS (equivalence holds for arbitrary valuations): The APS is the value of worst bundle in the best b_i- balanced distribution.

For equal entitlements ($b_i = \frac{1}{n}$ for integer n), the maximin share MMS definition requires a balanced distribution over disjoint bundles.

The APS is a relaxed version of the MMS (bundles need not be disjoint, $APS \geq MMS$) that extends to arbitrary entitlements.

A homogeneous divisible good

m identical items.

m is sufficiently large so that $b_i \cdot m$ is integer for all agents i.

$APS_i = v_i(b_i \cdot m)$.

- Upper bound: price each item as $\frac{1}{m}$.
- Lower bound: b_i- balanced distribution with m bundles, each of $b_i \cdot m$ consecutive items (in cyclic order) and probability $\frac{1}{m}$.

A homogeneous divisible good

$APS_i = v_i(b_i \cdot m)$.

- Lower bound: b_i- balanced distribution with m bundles, each of $b_i \cdot m$ consecutive items (in cyclic order) and probability $\frac{1}{m}$.

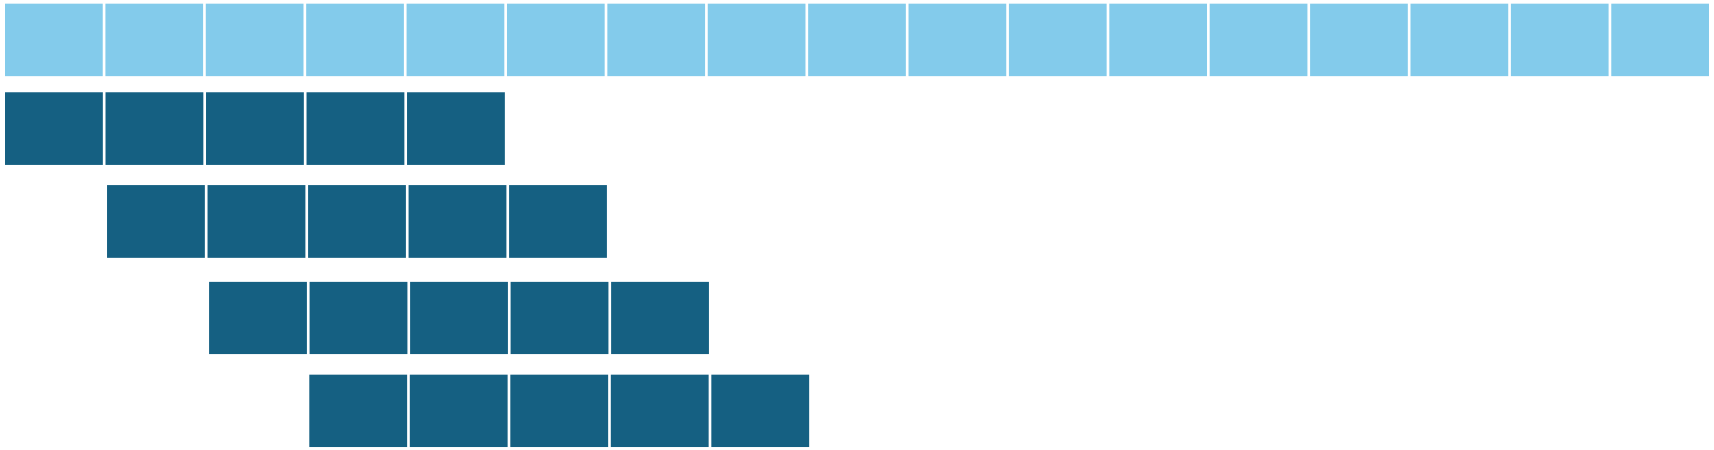

A homogeneous divisible good

When dividing a homogeneous divisible good, the proportional allocation is an APS allocation, regardless of the valuations of the agents.

Moreover, if valuations are strictly monotone, then the proportional allocation is the only APS allocation.

The APS is not a feasible share

There are instances with indivisible items, even with equal entitlement and identical additive valuations, in which in every allocation, some agent does not get her APS.

(In these instances, $APS > MMS$)

We settle for ρ-APS allocations, in which every agent gets at least a ρ fraction of her APS.

Ex-ante share-based fairness notions

Maximum expectation share MES.

Recall: a distribution is b_i-balanced if every item has probability b_i of appearing in a random bundle.

The APS is the value of worst bundle in the best b_i-balanced distribution.

The MES is the expected value of a random bundle in the best b_i-balanced distribution.

MES is at least as large as APS, and at least as large as the proportional share.

A dominating ex-ante share $\frac{1}{k+1} < b_i \leq \frac{1}{k}$

Recall: The MES is the expected value of a random bundle in the best b_i-balanced distribution.

For $\frac{1}{k+1} < b_i \leq \frac{1}{k}$, let $\widehat{b_i} = \frac{1}{k}$.

The $\widehat{MES}$ is the expected value of a random bundle in the best $\widehat{b_i}$-balanced distribution.

$\widehat{MES}$ is an upper bound on every feasible ex-ante share, and a tight upper bound for additive valuations.

Summary for share based notions $\frac{1}{k+1} < b_i \leq \frac{1}{k}$

Upper bounds on all feasible shares, tight for additive valuations.

Ex-post: $\widehat{MMS}(v_i, b_i) = MMS(v_i, \frac{1}{k})$.

Ex-ante: $\widehat{MES}(v_i, b_i) = MES(v_i, \frac{1}{k})$.

The anyprice share (APS): a share that induces the proportional allocation for a divisible homogeneous good (such as money).

The proportional share: a share that is always feasible ex-ante.

Approximations for additive valuations

Ex post:

For which values of ρ is ρ-APS feasible?

For which values of ρ is ρ-$\widehat{MMS}$ feasible?

Ex ante:

MES is feasible. For which values of ρ is ρ-$\widehat{MES}$ feasible?

BoBW?

Ex-ante approximation of $\widehat{MES}$

Recall, MES is feasible ex-ante, agent i gets M with probability b_i.

Problem for $\widehat{MES}$: after rounding up, sum of entitlements exceeds 1.

For $\frac{1}{k+1} < b_i \leq \frac{1}{k}$, let $\widehat{b_i} = \frac{1}{k}$.

Proposition: if $\sum_i b_i = 1$ then $\sum_i \widehat{b_i} < \gamma \simeq 1.69103$. Moreover, γ is the smallest possible universal upper bound.

ρ-$\widehat{MES}$ for $\rho = \frac{1}{\gamma} \simeq 0.591$: agent i gets M with probability $\frac{\widehat{b_i}}{\sum \widehat{b_j}} \geq \frac{1}{\gamma} \cdot \widehat{b_i}$.

Ex-post approximation of $\widehat{MMS}$

Approximation ratio is no better than $\frac{1}{\gamma} \simeq 0.591$.

In fact, approximation ratio is no better than 0.5.

$m = 2n - 2$ items of value 1.

$n - 1$ agents of entitlement $\frac{1}{n} < b_i < \frac{1}{n-1}$.

Then $\widehat{b_i} = \frac{1}{n-1}$ and $\widehat{MMS} = 2$.

One agent of entitlement $\frac{1}{n+1} < b_n < \frac{1}{n}$. Then $\widehat{b_n} = \frac{1}{n}$ and $\widehat{MMS} = 1$.

In every allocation, at least one agent does not get more than $\frac{1}{2} \cdot \widehat{MMS}$.

Approximation of $\widehat{MMS}$

Theorem [Babaioff and Feige, 2024]:

- For every allocation instance with arbitrary entitlements and additive valuations, there is a $\frac{1}{2}$-$\widehat{MMS}$ allocation.
- For every $\varepsilon > 0$, a $(\frac{1}{2} - \varepsilon)$-$\widehat{MMS}$ allocation can be computed in polynomial time.
- The results hold also for $\widehat{TPS}$ instead of $\widehat{MMS}$.

For additive valuations, the truncated proportional share (TPS) is a share whose value is sandwiched between the MMS and the proportional share. Moreover, unlike the MMS, its value can be computed in polynomial time.

Proof technique

Consider the bidding game, an allocation mechanism for agents of arbitrary entitlements.

Design a safe strategy for an agent playing the bidding game, that guarantees at least $\frac{1}{2} \cdot \widehat{MMS}$, regardless of the strategies used by other agents.

Existence of $\frac{1}{2} \cdot \widehat{MMS}$ allocations follows, as all agents can use the safe strategy.

If the safe strategy can be implemented in polynomial time, then a $\frac{1}{2} \cdot \widehat{MMS}$ allocation can be found in polynomial time.

The bidding game [Babaioff, Ezra, Feige 2021] (a *poorman* game [Lazarus, Loeb, Prop, Ullman 1995])

The bidding game for allocating m items takes m rounds.

Initially, each agent i gets a budget equal to her entitlement b_i.

In every round:

- Each agent makes a bid not higher than her remaining budget.
- Highest bidder wins (ties broken at random), pays her bid, and picks an item of her choice.

Game ends when no items are left (or no budget is left).

Warm-up: the proportional bidding strategy

Scale the additive valuation function v_i so that $v_i(M) = 1$.

Assume (for simplicity) that no single item has value above b_i.

A safe strategy that achieves at least $\frac{1}{2} \cdot PS$

(without the assumption, can guarantee $\frac{1}{2} \cdot TPS$).

- In every round, bid value of highest remaining item. If insufficient budget remains, then bid the full remaining budget (an under-bid).
- If the bid wins, take highest value item and pay your bid.

Bidding against an adversary

- $\sum_i b_i = 1 \; and \; v(M) = 1.$
- An adversary aims to minimize the bidder value.
- Proportional strategy:
  bid *min(highest value item, budget)*

fraction of PS achieved $\geq$ fraction of budget spent

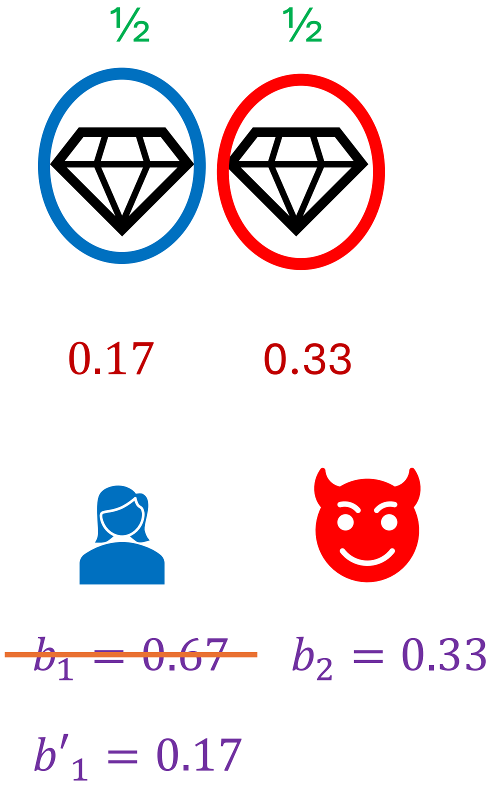

$\frac{1}{2}PS$ guarantee for proportional bidding strategy

If agent i spends at least $\frac{b_i}{2}$, she gets value at least $\frac{b_i}{2}$, as desired.

Assume for the sake of contradiction that agent i managed to spend less than $\frac{b_i}{2}$.

This implies that i never needed to under-bid.

Other agents must pay at least agent i's bids for the items that they win.

Regardless of which items they choose, they consume value at most $1 - b_i$ (their total budget).

A value of $\frac{b_i}{2}$ still remains, contradicting the assumption.

The proportional strategy ensures at least $\frac{1}{2}TPS$.

We need to achieve $\frac{1}{2}\widehat{MMS}$, and in fact achieve $\frac{1}{2}\widehat{TPS}$.

The proportional strategy is too weak to achieve this.

Bidding against an adversary

- $\sum_i b_i = 1 \; and \; v(M) = 1.$
- An adversary aims to minimize the bidder value.
- Proportional strategy:
  bid *min(highest value item, budget)*

- Is there a better strategy?
- Yes!

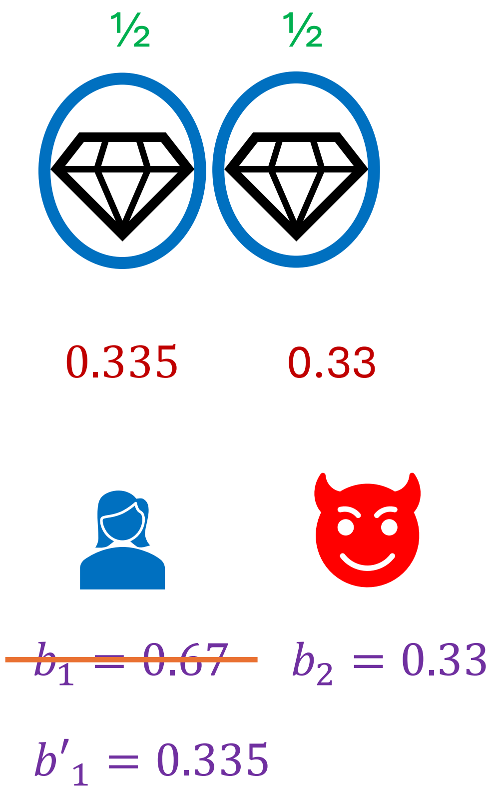

Safe strategy for $\frac{1}{2}\widehat{MMS}$

Unlike the proportional strategy which ensures $\frac{1}{2}TPS$, any strategy that achieves $\frac{1}{2}\widehat{MMS}$ (and $\frac{1}{2}\widehat{TPS}$) must be:

- Non myopic (the bid depends on values of multiple items).
- Non monotone (bids may increase across rounds).

Designing a safe strategy for $\frac{1}{2}\widehat{MMS}$ requires a much better understanding of the strategy space for the bidding game.

Understanding the bidding game

- Safe strategies: guarantee a value $V_i(b_i)$ for b_i, for any adversary.

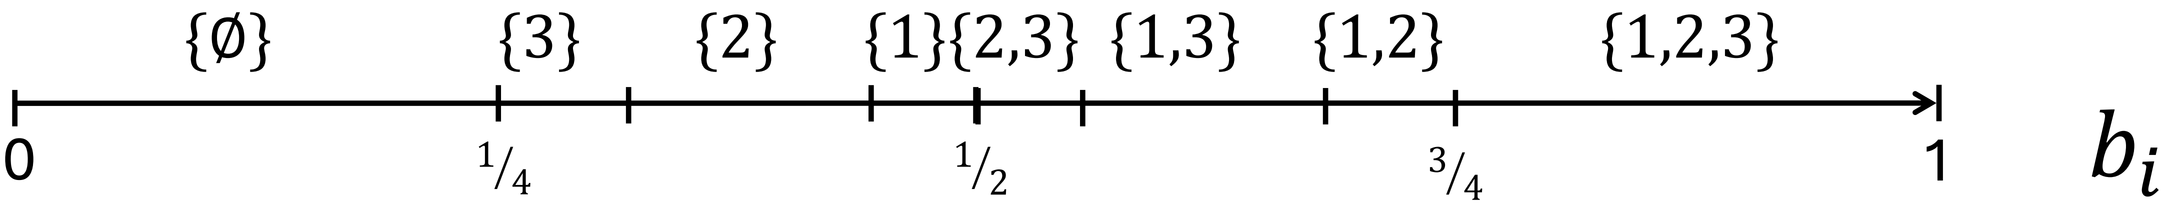

- All subsets, by increasing order of values.
- Symmetry: $V_i(b_i) + V_i(1 - b_i) = V_i(M)$ for non-transition b_i.
- Entitlement > ½ ➔ at least ½ of $v_i(M)$.
- Entitlement = ½ ➔ at least MMS of two agents.
- Non-constructive!
- NP-hard to compute a worst-case-optimal safe strategy.

Safe strategy for $\frac{1}{2}\widehat{MMS}$

Our strategy is explicit and polynomial time, as long as the agent has at most half of the remaining budget.

If at any point the agent has more than half of the remaining budget, the strategy switches to the non-constructive existential result.

For this last reason, our strategy is not known to be implementable in polynomial time.

For every ε, we have a polynomial time version that achieves $(\frac{1}{2} - \varepsilon)\widehat{TPS}$.

Additional notes on the bidding game

There are several natural variations on this game (pay second price, choose multiple items).

The strongest currently known approximations for APS allocations were proved by designing safe strategies for the bidding game.

$\frac{3}{5}$-APS for additive valuations [Babaioff, Ezra, Feige 2021].

$\frac{1}{3}$-APS for submodular valuations [Ben Uziahu and Feige 2023].

(Submodular is the discrete analog of concave valuations, non-increasing marginal values.)

Best of Both Worlds Impossibility

Cannot ensure constant fraction of $\widehat{MMS}$ ex-post and constant fraction of $\widehat{MES}$ ex-ante simultaneously.

Example: two identical items and three agents with entitlements (0.34,0.34,0.32).

- Each of agent 1 and agent 2 must get an item in the ex-post allocation. Otherwise, no approximation of the $\widehat{MMS}$.
- Ex-ante, agent 3 must have positive probability of receiving an item. Otherwise, no approximation of the $\widehat{MES}$.

These two requirements are incompatible.

Summary for additive valuations

Ex-post, can ensure every agent at least $\frac{1}{2}\widehat{MMS}$, and hence at least half of every feasible share. This is best possible.

Ex-ante, can ensure every agent at least $0.591\widehat{MES}$, and hence at least 0.591 of every feasible share. This is best possible.

No BoBW: every randomized allocation that gives every agent positive value ex-ante, must have in its support an allocation in which some agent gets 0-value (ex-post), despite there being a feasible share of positive value for the agent.

Chores

- Chores have positive costs (negative values).
  - Doing the laundry
  - Cleaning the bathroom
  - Sweeping the floors
- All chores must be assigned.
- Agents have responsibilities (sum to 1).
- Agents aim to minimize their costs.
- A share bounds the maximal cost the agent might suffer.

Share based fairness definitions for chores

Maximin (MMS) changes to minimax (an n-partition that minimizes the cost of the most costly bundle).

Ex-post dominating share: $\widetilde{MMS}$ (responsibility $\frac{1}{k+1} \leq b_i < \frac{1}{k}$ rounded down to $\frac{1}{k+1}$).

Ex-ante dominating share: $\widetilde{MES}$.

BoBW result for chores

Theorem: There is a polynomial time randomized assignment algorithm that for every input instance with additive cost functions and arbitrary responsibilities outputs a distribution over assignments for which for every agent:

- Ex-ante: the expected cost is at most her proportional share (hence, at most 2-$\widetilde{MES}$).
- Ex-post: the cost is at most 2-$\widetilde{MMS}$.

Moreover, the approximation ratio of 2 is best possible, both ex-ante and ex-post.

A general technique for additive valuations: faithful implementation of fractional allocations

Find a fractional allocation A^* with good fairness properties.

(E.g., each agent i gets a b_i fraction of each item.)

- Faithful rounding: a rounded integral solution A in which every agent gets at least her fractional value, up to one item. Provides an ex-post guarantee.
- Faithful implementation: a distribution over faithful integral solutions, where the expectation of the distribution is exactly A^*. Adds an ex-ante guarantee - every agent gets her fractional value in expectation.

Faithful implementation gives “best of both worlds”: ex-ante and ex-post.

Some relevant work on faithful implementations

[Srinivasan 2008] – designed and applied a faithful implementation (*with spread of at most one item*) so as to get a $\frac{3}{4}$ approximation to the maximum welfare with budget additive valuations.

[Budish, Che, Kojima and Milgrom 2013] – motivated the concept of faithful implementations, and generalized it to bi-hierarchies.

[Freeman, Shah and Vaish 2020] – presented two “best of both worlds” results for fair allocation.

[Aziz 2020] – a simplified randomized allocation that is ex-ante envy free and ex-post envy free up to one good.

The unifying lemma

Lemma. Every fractional allocation A^* can be implemented in polynomial time as a distribution over (polynomially many) integral faithful allocations A^1, A^2 … .

Implementation: $A^* = \sum_k \lambda_k A^k$, where $\sum_k \lambda_k = 1$, and $\lambda_k \geq 0$ for all k.

Every allocation A^k is faithful: every agent i receives same value as in A^*, up to one item.

We shall present the proof given in [Aziz 2020].

Birkhoff – von Neumann theorem [BvN]

- Every doubly stochastic matrix decomposes (in polynomial time) into a convex combination (distribution) over permutation matrices.

0.6	0.2	0.2
0.3	0.1	0.6
0.1	0.7	0.3

		0.1
	0.1	
0.1		

	0.2	
0.2		
		0.2

		0.1
0.1		
	0.1	

0.6		
		0.6
	0.6	

A BvN algorithm for faithful implementation

Consider a fractional allocation. For example:

	Agent 1	Agent 2
Item 1	0.4	0.6
Item 2	0.3	0.7
Item 3	1	0

A BvN algorithm for faithful implementation

Allocate the integral items (item 3 to agent 1):

	Agent 1	Agent 2
Item 1	0.4	0.6
Item 2	0.3	0.7
Item 3	1	0

	Agent 1	Agent 2
Item 1	0.4	0.6
Item 2	0.3	0.7

A BvN algorithm for faithful implementation

Add dummy items (of value 0) to get integer column sums:

	Agent 1	Agent 2
Item 1	0.4	0.6
Item 2	0.3	0.7

	Agent 1	Agent 2
Item 1	0.4	0.6
Item 2	0.3	0.7
Item 4 (dummy)	0.3	0.7

A BvN algorithm for faithful implementation

	Agent 1	Agent 2
Item 1	0.4	0.6
Item 2	0.3	0.7
Item 4 (dummy)	0.3	0.7

Break agents into clones to get column sum 1.

Clones “eat” fractions in order of value of items.

	Agent 1	Agent 2a	Agent 2b
Item 1	0.4	0.3	0.3
Item 2	0.3	0.7	0
Item 4 (dummy)	0.3	0	0.7

Now we have a doubly stochastic allocation matrix.

How did we partition the fractional allocation of Agent 2 among his clones?

Recall, fractional allocation for Agent 2:
0.6 (item 1) + 0.7 (item 2) + 0.7 (item 4)

Sort items by value to Agent 2. Suppose that $v_2(2) \geq v_2(1) \geq v_2(4)$.
0.7 (item 2) + 0.6 (item 1) + 0.7 (item 4)

Break fractional allocation so that partial sums are integers:
0.7 (item 2) + 0.3 (item 1) + 0.3 (item 1) + 0.7 (item 4)

How did we partition the fractional allocation of Agent 2 among his clones?

Recall, fractional allocation for Agent 2:
0.6 (item 1) + 0.7 (item 2) + 0.7 (item 4)

Sort items by value to Agent 2. Suppose that $v_2(2) \geq v_2(1) \geq v_2(4)$.
0.7 (item 2) + 0.6 (item 1) + 0.7 (item 4)

Break fractional allocation so that partial sums are integers.
0.7 (item 2) + 0.3 (item 1) + 0.3 (item 1) + 0.7 (item 4)

Clone 2a Clone 2b

A BvN algorithm for faithful implementation

Implement the doubly stochastic allocation as a distribution over permutation matrices.

	Agent 1	Agent 2a	Agent 2b
Item 1	0.4	0.3	0.3
Item 2	0.3	0.7	0
Item 4 (dummy)	0.3	0	0.7

$$\begin{bmatrix} 0.4 & 0.3 & 0.3 \\ 0.3 & 0.7 & 0 \\ 0.3 & 0 & 0.7 \end{bmatrix} = \begin{bmatrix} & & 0.3 \\ & 0.3 & \\ 0.3 & & \end{bmatrix} + \begin{bmatrix} 0.4 & & \\ & 0.4 & \\ & & 0.4 \end{bmatrix} + \begin{bmatrix} & 0.3 & \\ 0.3 & & \\ & & 0.3 \end{bmatrix}$$

In every permutation matrix, every clone gets exactly one item.

In the corresponding allocation, every agent gets all items that are allocated to his clones. (Agent 2 has two clones - gets two items.)

Discard dummy items. (Discard item 4.)

The BvN algorithm produces an implementation

Ex-ante: In expectation, every agent i gets every item j with probability exactly equal to the fractional allocation of j to i.

- The fractional allocation of an agent is split among her clones.
- For every clone, the probability of getting an item j is exactly equal to the fraction of j assigned to the clone. (BvN).
- Every agent gets all items of her clones.

The implementation is faithful

Ex-post: Every agent gets same value as she does in the fractional allocation, up to one item.

Even stronger: The best and the worst ex-post allocation to agent i have the same value, up to one item (the spread is at most one item).

- Removing best item from best possible allocation B_i, we get an allocation that is no better than worst possible allocation W_i.

0.7 (item 2) + 0.3 (item 1) + 0.3 (item 1) + 0.7 (item 4)

Clone 2a Clone 2b

$$v_i(\max c_{2a}) \geq v_i(\min c_{2a}) \geq v_i(\max c_{2b}) \geq v_i(\min c_{2b}) \geq \cdots$$

Summary: the faithful implementation lemma

Lemma. Every fractional allocation A^* can be implemented in polynomial time as a distribution over (polynomially many) integral faithful allocations A^1, A^2 ... .

Implementation: $A^* = \sum_k \lambda_k A^k$, where $\sum_k \lambda_k = 1$, and $\lambda_k \geq 0$ for all k.

Every allocation A^k is faithful: every agent i receives same value as in A^*, up to one item. (Moreover, the spread is at most one item.)

The lemma has many applications.

Allocation of chores, additive valuations

A faithful implementation of the proportional fractional allocation (each agent i gets a b_i fraction of each item).

Ex-ante: each agent i gets each item with probability b_i. For additive valuations, get in expectation the proportional cost $b_i \cdot v_i(M)$. This is at most $2\widetilde{MES}$ (worst case is responsibility $1-\varepsilon$).

Ex-post: each agent i with entitlement $\frac{1}{k+1} \leq b_i < \frac{1}{k}$ gets value at most $v_i(e_1) + \frac{1}{k} v_i(M \setminus \{e_1\}) = \frac{k-1}{k} v_i(e_1) + \frac{1}{k} v_i(M)$

$$\leq \frac{k-1}{k}\widetilde{MMS} + \frac{k+1}{k}\widetilde{MMS} = 2\widetilde{MMS}.$$

Summary

When agents have arbitrary entitlements and are assumed to have additive valuations, which allocation mechanisms can we recommend?

Goods: the bidding game is a plausible candidate. Agents have safe strategies that guarantee at least half of every feasible share. This is a best possible worst-case guarantee ex-post.

Chores: a faithful implementation of the proportional allocation. Best worst-case ex-ante and ex-post guarantees simultaneously.

Some open questions

Propose a useful version of envy freeness for arbitrary entitlements.

Additive valuations:

- Do $(1-\varepsilon)$-MMS allocations become feasible as n grows?
- Do APS (or even just $\widetilde{MMS}$) allocations exist whenever entitlements sum up to less than 1?

Are there $\frac{1}{2}$-MMS allocations for subadditive valuations?